ELSEVIER

Fusion Engineering and Design 32–33 (1996) 323–335

Fusion Engineering and Design

Physics design and scaling of Elise

E.P. Lee[a], R.O. Bangerter[a], C.F. Chan[a], A. Faltens[a], J. Kwan[a], E. Henestroza[a], K. Hahn[a], P. Seidl[a], J.J. Barnard[b], A. Friedman[b], D.P. Grote[b], W.M. Sharp[b]

[a]*Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA*
[b]*Lawrence Livermore National Laboratory, Livermore, CA 94550, USA*

Abstract

Elise is an electrostatically focused heavy-ion accelerator being designed and constructed at Lawrence Berkeley Laboratory. The machine is intended to be the first half of the four-beam Induction Linear Accelerator (LINAC) Systems Experiment (ILSE), which will ultimately test the principal beam dynamics issues and manipulations of induction heavy-ion drivers for inertial fusion. Elise will use an existing 2 MeV injector and will accelerate space-charge-dominated pulses to greater than 5 MeV. The design objective of Elise is to maximize the output beam energy within the fixed project budget, while allowing for adequate beam diagnostics, flexibility in the acceleration schedule, and beam parameters suitable for ILSE and the experimental program. We review the design equations and 'rules of thumb' used for choosing the beam and lattice parameters for heavy-ion induction accelerators, and we discuss incorporating these relations in a spreadsheet program that generates internally consistent lattice layouts and acceleration schedules. These designs have been tested using a one-dimensional particle simulation code (SLIDE), a three-dimensional fluid/envelope code (CIRCE) and a three-dimensional particle-in-cell code (WARP3d). Sample results from these calculations are presented. Results from these dynamics codes are also shown, illustrating sensitivities to beam and lattice errors, and testing various strategies for longitudinal confinement of the beam ends.

1. Introduction

Elise is an electrostatically focused heavy-ion induction accelerator being designed and constructed at Lawrence Berkeley National Laboratory (LBNL). The machine will be built in a way that allows future expansion into the full Induction Linear Accelerator (LINAC) Systems Experiment (ILSE) configuration [1] that includes beam-combining and magnetic focusing. Elise will demonstrate beam manipulations and address other beam dynamics issues associated with a heavy-ion driver for inertial fusion.

The design objective of Elise is to optimize the beam current density and the total energy, while allowing for adequate beam diagnostics, flexibility in the acceleration schedule, and beam parameters suitable for ILSE and other experimental programs. Since Elise has many features similar to those of driver-scale accelerators, the project will produce valuable information (both technical and cost data) for designing a fusion driver.

Elise will have an array of four electrostatic quadrupole (ESQ) focusing channels capable of transporting up to (and, for some scenarios, more than) a total of 3.2 A of beam current. With a line charge density similar to that of heavy-ion drivers,

PII S0920-3796(96)00485-1

Elise will accelerate a beam pulse (duration of 1 μs) of K^+ ions from an initial energy of 2 MeV to a final energy of more than 5 MeV. Initially, only one beam channel will be used during nominal Elise operation. Eventually, for ILSE upgrading, a new injector will produce four beams that will be combined into one before entering a magnetically focused accelerating section for further acceleration to more than 10 MeV.

Elise is intended to approximate as realistically as possible the front end of an induction LINAC driver. The electric quadrupole geometry and operating voltages, and the resultant transportable beam are all at full driver scale; except for the necessarily shorter pulse length they would provide directly useful information for the future. In a driver, the electric focusing section would be a few hundred meters in length, throughout most of which the beam pulse is almost of constant physical length and constant line charge density. Part of Elise will duplicate such operation. At the entrance to the driver, some tens of meters are devoted to the process of launching the bunch in a manner that avoids heating it in the longitudinal phase plane; this space is insignificant on the driver scale. For Elise, a gentle starting section would correspond to a significant fraction of the entire machine, yet it is important to devise starting scenarios for a driver that can be investigated experimentally in Elise. Consequently the machine is being designed to enable the exploration of several candidate scenarios, all of which are acceptable in preserving beam quality, but will have various economic consequences.

2. Elise technical design parameters

Approximate design and performance goal parameters are listed in Table 1. In some cases (such as final current), these exceed commissioning requirements [2]. They will also vary depending on the mode of operation and experimental objectives.

Table 1
Elise technical design parameters

Parameter	Units	Value
Initial ion kinetic energy	Megaelectronvolts	2
Initial beam current	Amperes	0.8
Initial pulse duration	Microseconds	1.0
Initial beam line charge density	Microcoulombs per meter per beam	0.25
Initial number of beams	None	1
Final average ion kinetic energy	Megaelectronvolts	5
Final beam current	Amperes	1.0
Final pulse duration	Microseconds	0.8
Final beam line charge density	Microcoulombs per meter per beam	0.2
Final beam energy	Joules	4
Length of LINAC	Meters	15
Number of acceleration gaps	None	51
Ion mass number	Atomic mass units	39
Ion charge	Electron charge	+1

3. ESQ injector and matching section

On the basis of reliability, driver scalability and beam specifications, an ESQ [3] design was selected for the injector. A one-beam prototype of the ESQ injector has been built and tested at LBNL. The prototype provides 0.8 A or more of 2 MeV K^+ ions—equivalent to a line charge density of 0.25 μC m^{-1}—at a low normalized emittance (1 mm mrad), a repetition rate of 1 Hz and a pulse length of 1 μs. The ESQ injector consists of a diode, followed by a sequence of quadrupoles arranged to focus and accelerate the beam at the same time. A schematic diagram of the injector is shown in Fig. 1. The ESQ is

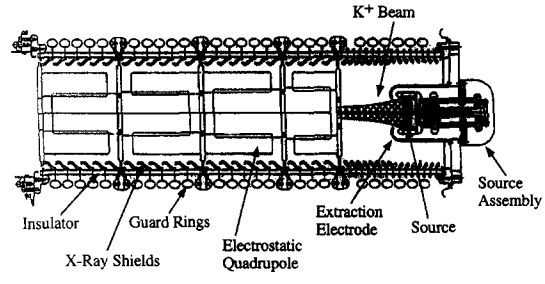

Fig. 1. Schematic diagram of ESQ injector.

generally a long machine, with correspondingly low gradients. The secondary electrons are swept out by the large transverse fields, which reduces significantly the breakdown risks. In addition, the sources in an ESQ are generally smaller than those matched to a high gradient column, so the intrinsic emittance is lower. The ESQ is also attractive because of its scaling; it has the potential advantage of operating at energies somewhat higher than 2 MeV, since the critical issues of emittance growth in an ESQ tend to center on the transition from the diode to the first accelerating quadrupoles.

The diode consists of a hot alumino-silicate source with a large curved emitting surface surrounded by a thick 'extraction electrode'. An extraction pulser switches the source from -80 kV to $+80$ kV relative to the extraction electrode during beam turn-on. The pulser voltage profile has a 0.5 μs rise-and-fall time, with a 1.0 μs flat-top. The EGUN [4] code was used in the design of the diode.

The design of the ESQ injector was based on the three-dimensional particle-in-cell (PIC) codes WARP3d [5] and ARGUS [6] running in a steady state mode. A full three-dimensional PIC simulation code was required to incorporate the beam space-charge field as well as the self-consistent fields from the accelerating quadrupoles, including their 'interdigital' structure (cantilevered electrode rods). The interdigital structure of the ESQs could enhance the aberrations. The resulting kinematic distortions lead to S-shaped phase spaces that if not corrected, will lead eventually to emittance growth. These beam aberrations can be minimized by increasing the injection energy and/or strengthening the beam focusing. The parameters of the design represent optimal choices to obtain a proper balance between breakdown risk and emittance growth. With the design diode voltage of 750 kV, a small normalized emittance (0.6 mm mrad) was obtained. The intrinsic normalized emittance is 0.4 mm mrad.

The experimental results from the injector matched the code predictions [7]. Measurements of the transverse phase space distribution have shown excellent agreement with WARP3d calculations. Good agreement was also found in the current profile with the transient longitudinal dynamics of the beam in the ESQ simulated by running GYMNOS [8] and WARP3d in time-dependent modes.

The matching Section (Fig. 2) was designed to reduce the beam size from the injector to the size of the main transport channel. Using WARP3d, a simulated beam was created in the injector and transported through the matching section. With the rise-and-fall time of 0.5 μs from the injection pulser, the rise and fall at the end of the matching section is between 0.3 and 0.5 μs, i.e. short enough to alleviate concerns about beam expansion in the matching section from the lack of axially confining fields.

4. ESQ transport dynamics in the induction accelerator

The transportable beam current in an ESQ channel is governed by the aperture radius and the voltages applied to the ESQ. The scaling is such that the optimum beam current per unit cross-sectional area is obtained by using an array of ESQ channels with a small aperture radius. However, there is a lower practical limit to the aperture radius, as a result of the minimum beam clearance imposed by random errors in beam steering and accelerator alignments. In the present design, Elise beam channels will have a constant aperture radius of 2.33 cm and a nearly constant quadrupole voltage of about ± 60 kV.

Based on optimizing the 'figure of merit', defined as the maximum beam line charge per unit area, the optimum aperture radius is determined to be $b = 2.33$ cm. The peak beam radius a is related to the aperture radius by

$$b = 1.25a + c$$

The coefficient 1.25 arises from a limitation from the image force from the electrodes, and the beam clearance c is obtained from an estimate of the beam steering random error and acclerator alignment limits, which are expected to be of the order of a few millimeters. The maximum line charge density could be more than doubled if the clearance allowance can be reduced from 1.0 cm to less

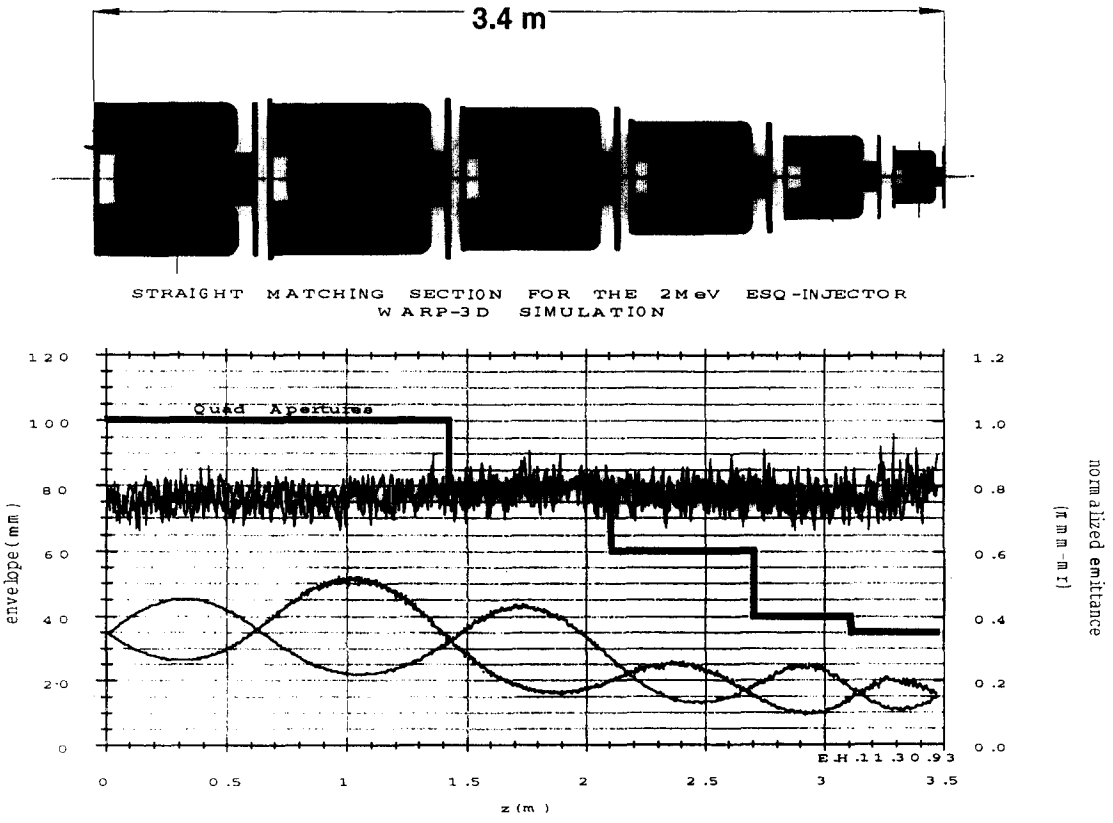

Fig. 2. Matching section for the Elise accelerator. As drawn here, the beam travels from left to right, as it is compressed and matched to the quadrupolar geometry.

than 0.2 cm. The electrode radius R_e is determined by setting the dodecapole component of the focusing electric field to zero, which means that the ratio of the electrode radius to the aperture radius is 1.146 (or about 8/7).

ESQ spark-down tests (at LBNL) have shown that the maximum voltages for the ESQ is proportional to the square root of the spacing between the quadrupole electrodes. An ESQ with $b = 2.3$ cm and $R_e = 2.53$ cm broke down with 230 kV between the quadrupole electrodes. As a rule of thumb, we set the normal operating point at below 50% of the spark-down threshold value. Based on these rules, we found that the optimum value of b is 2.33 cm and the ESQ voltage V_q is 118 kV (or ± 59 kV).

The next important dimension to determine is the lattice half-period length L and the maximum associated transportable line charge density λ. These are accurately determined from the matched envelope relations [9]

$$\cos\sigma_0 = 1 - \frac{\eta^2(1-2\eta/3)}{2}\left(\frac{E'}{2T}\right)^2 L^4 \frac{\lambda}{4\pi\varepsilon_0 T}\left(\frac{2L}{a}\right)^2 = 2(1-\cos\sigma_0) \times \left[1+\eta\left(\frac{E'}{2T}\right)\frac{L^2}{4}\left(1-\frac{\eta}{2}\right)\right]^{-2}$$

Here, σ_0 is the undepressed phase advance per lattice period, η is the field-occupying fraction of the ESQ, $E' = V_q/b^2$ is the quadrupole field gradient, and T is kinetic energy (expressed in electron-volts).

The occupancy factor η can be approximated according to

$$\eta = (L - 6.0\text{ cm})/L$$

the spacing of 6.0 cm is made up of the 'dead space'

between the ESQ electrodes and the end-plates (which is necessary for holding the quadrupole voltage), the end-plate thickness (about 1 cm), and an acceleration gap of 2.0 cm between the end-plates of neighboring ESQ modules (for typical 100 kV acceleration). Obviously, η grows with L; typical values of η start from about 0.7 at the beginning and exceed 0.8 at the end of Elise.

According to the above equation for the half-lattice period L, and using a 75° undepressed phase (85° is considered the upper limit of stable operation), we found L to be 20.8 cm at $T=2$ MeV. With everything else being kept constant, L is proportional to the square root of the beam energy (a reasonable approximation for slowly varying η). However, once the dimensions, such as b and L, are fixed, the undepressed phase advance can be raised up to 85° if the quadrupole voltage V_q can reach ± 65.5 kV without incurring spark-downs.

Ideally, to avoid a mismatch of the beam radius, the half-lattice period L should vary as the square root of the beam energy, and inversely as η. Thus, L becomes longer after each acceleration gap. Depending on the cost of engineering and fabricating ESQs, it may be worthwhile to consider using blocks of ESQs that have the same L within each block, and accept the consequence of a minor mismatch. The net result is a reduction in the transportable beam current compared with the case of continuously changing L in each period. For example, assuming no beam head-to-tail energy variation, a single-stage L will transport 50%, a three-stage L will have 73% and a 10-stage L will have 93% of the beam current found in the continuous limit. (The advantage of a 10-stage or continuous limit may not be this high if the acceleration schedule imposes a significant head-to-tail energy variation.) At present, the estimated cost of fabricating ESQs with continuously varying L is about US$350 000 (35%) higher than the cost of ESQs with only two different L values. Thus, Elise will most likely use continuouly varying L.

Voltages on the ESQ electrodes can be arranged either in the bipolar or unipolar configurations. For bipolar arrangement potentials of ± 59 kV are applied to the two pairs of electrodes, whereas $+118$ kV and ground potentials are used in unipolar configurations. The two cases produce a small difference in the beam dynamics, as long as the beam energy is much higher than the focusing potential. Basically, the unipolar design has twice the acceleration voltage across an accelerating gap, but only half as many accelerating gaps (there is a non-accelerating gap between each pair of accelerating gaps). The emittance growth that results from the larger kicks is undesirable but tolerable in the unipolar case. In the end, the selection between bipolar and unipolar quadrupoles relies mostly on engineering and economics reasons.

Since the estimated costs for both systems are the same (within 5%), we have chosen the bipolar design, because of its lower risk and higher flexibility. The ability to use cantilevered mounting of the ESQs makes it possible to achieve a reasonable axial packing of core material.

To prevent beam elongation, a small head-to-tail velocity tilt is applied, i.e. the tail is moving faster than the head, so the tune at the beam head is higher than that of the tail, and the tail beam is 'fatter' than the head beam. For a constant aperture but aperiodic transport channel, it is appropriate to define matching to mean constant maximum envelope radius. The gentle deviation from the matched envelope that results from velocity tilt is unimportant, as shown in simulations, as long as the envelope stays in the dynamic aperture mentioned above. Envelope codes predict additional ± 2.0 mm of mismatch oscillation at the beam ends, as a result of the velocity tilt.

It may be desirable to have the same voltages on all the quadrupoles, so that several quadrupoles are driven by a common power supply. For a given acceleration schedule, it is straightforward to design a lattice that satisfies the above criteria.

5. Longitudinal dynamics and acceleration schedule

Assuming that the transverse beam dynamics are handled properly by the appropriate ESQ focusing field and the correctly scaled half-length periods, we now examine the problem of longitudinal beam dynamics.

The acceleration schedule in Elise is closely tied to the lattice design. As discussed elsewhere [2] and in Section 4 of this paper, the present design has the length of quadrupoles increasing along the lattice approximately in proportion with the beam velocity, while keeping the same voltage on all quadrupoles. This choice avoids transverse mismatches, by maintaining a constant average focusing strength. Although the beam velocity as a function of position is, in principle, quite flexible, the choice made in Elise is to maximize the packing density of acceleration cores, consistent with breakdown limits, access for pulsed-power feeds and diagnostics, and preservation of some modularity. The juggling of parameters to optimize cost and performance is handled by the MATHEMATICA optimization code discussed in Section 6. The present design has average acceleration voltages that vary somewhat erratically between 55.7 and 153.75 kV, with gaps for diagnostics in place of two cells. Despite the irregular acceleration, simulations using the envelope/fluid code CIRCE indicate that the beam remains near transverse equilibrium along its length.

For any acceleration schedule, there remains considerable flexibility in the longitudinal compression of the beam. CIRCE simulations indicate that compression is best achieved by imposing a smooth velocity variation or 'tilt' along the beam, while balancing the beam axial space-charge field on the average by 'ear' fields applied in selected acceleration gaps. 'Ears' are small voltages superimposed on the original voltage at the beginning and end of the pulse. They are produced by fast-responding cores and pulesers. A velocity tilt is introduced by appropriate time variation of the accelerating voltage in one or more gaps, and theoretical work by Kim and Smith [10] indicates that the waveform that gives self-similar compression of the beam current has a nearly linear voltage variation with time. The amount of velocity tilt that can be imposed is limited by three considerations as follows.

(1) The undepressed phase advance per lattice period σ_0 is approximately proportional to the inverse square of the beam velocity for electric focusing, so the requirement that σ_0 at the beam head must remain less than 90° either limits the velocity tilt or forces a reduction in the average focusing strength, driving up the beam size.
(2) Any velocity tilt complicates beam steering, because the transverse kick from electric or magnetic dipoles depends on velocity.
(3) The maximum acceleration field, usually at the beam tail, plus any ear field there, which is also maximum at the tail, cannot exceed the breakdown limit in any cell.

The first constraint is the most salient for Elise, typically limiting velocity tilt to about 20%. The second constraint is also important, because the Elise beam will ultimately be merged with three others in the beam-combiner of the ILSE. Since the combiner uses electric dipoles to bend the four beams, its velocity acceptance is limited to a few per cent at the final Elise energy. The third constraint is typically more important as a limit on the spacing between ear cells than on the allowable velocity tilt.

Elise is being designed with sufficiently large induction cores to allow several compression scenarios to be tested. The simplest choice is to impose no velocity tilt, in which case the Elise beam would have a constant duration in the absence of axial space charge; with space charge, the beam would increase in duration by about 15%. Since the average beam energy increases from 2 to 5.7 MeV, the beam length nearly doubles in this case. A simple prescription for imposing a substantial tilt is to apply triangular acceleration pulses, with zero voltage at the head and the nominal cell voltage at the tail. The maximum velocity difference along the beam is then controlled by the number of cells with relatively small time-varying voltage; it gives the largest peak velocity tilt for a given compression and, therefore, limits the allowable focusing strength.

An alternative method for imposing tilt is to apply trapezoidally varying voltages in all cells. By controlling the rate of voltage variation, the rate of compression can be adjusted so that the beam current remains near the allowable maxi-

mum [9], making optimum use of the accelerator aperture. The two drawbacks of this approach are the need to apply time-varying voltages at all cells and the sizable tilt that remains at the end of the lattice.

A final method of generating a velocity tilt is referred to as the 'load-and-go' schedule. In this approach, the entire beam is allowed to drift into the accelerator before the pulsed power is turned on, so that the beam head traverses fewer active cells than does the tail. If the cell voltages are turned on instantly, then this technique would introduce stepwise velocity changes along the beam, so initiating sizable space-charge waves. However, the cell rise time is typically longer than the beam transit time between cells and, in this case, the velocity variation imparted to the beam is acceptably smooth, provided that all cells have the same flat-top voltage. The load-and-go schedule has the distinct advantage of requiring only flat waveforms, but there are several deficiencies as well. Since the head is not accelerated until the entire beam is in the accelerator, the average beam energy is lower than for other schedules, although the difference would be negligible for a driver. Also, there is no independent control of the tilt and the beam velocity, and there is a somewhat higher level of high-frequency space-charge wave imposed on the beam than in the other scenarios. We expect to try all these compression schedules on Elise.

For all the schedules with significant beam compression, the final velocity tilt exceeds the acceptance of the ILSE beam-combiner. This residual tilt can be removed by imposing a backward tilt on the waveforms in the last few cells, provided that the velocity variation along the beam is approximately linear.

The charge distribution along the length of the beam bunch generates a longitudinal electric field that, in the long wavelength approximation, is given by

$$E_z = -\frac{g}{4\pi\varepsilon_0}\frac{\partial\lambda}{\partial z}$$

where λ is the line charge density and g is a geometrical factor that, for cylindrical geometry, is given (near the bunch ends) approximately by

$$g = \frac{1}{2} + 2\ln\left(\frac{b}{a}\right)$$

Here, b and a are the beam pipe and beam radii respectively. More accurate treatments of this topic are given in Refs. [11,12]. This longitudinal electric field modifies the shape of the beam pulse and induces energy variation at the head and tail, as it travels through the focusing quadrupoles and acceleration stations.

The voltage waveforms at the acceleration stations can be tailored to help control the beam pulse, in addition to the role of performing the acceleration. We consider here two examples of simulations of the longitudinal beam dynamics with the one-dimensional PIC code SLIDE. We assume that the incoming beam from the injector is adequately represented by a simplified pulse of trapezoidal shape. It has a rise time of 0.3 μs, a flat-top portion of 1.5 μs and a fall time of 0.4 μs. The beam energy is 2 MeV and $\lambda = 0.25$ μC m^{-1} ($I = 0.8$ A) at the flat-top.

The first example corresponds to the constant-current scenario. The beam is accelerated by a series of square voltage pulses. The pulse shape of the outgoing beam, at the end of the 54 half-length period of the quadrupoles, is shown in Fig. 3. On comparing with the incoming pulse shape (not shown), there is a decrease of about 10% in the length of the flat-top portion. The average beam energy is 5.7 MeV. There is also a small increase in energy (about 0.38%) at the rise portion and a small decrease (by about 0.55%) at the fall portion of the beam pulse. This is because of the effect of space charge in the bunch ends and is correctable by the application of voltage ears.

In the second example, the beam is accelerated and compressed by a series of trapezoidal voltages pulses, with a lower voltage for the head and higher voltage for the tail. The shape of the outgoing beam pulse, with a slant top, is shown in Fig. 4. Total pulse length is about 1.8 μs (down from 2.2 μs), with an average peak current 0.9 A. There is also a large energy tilt, with the head at 5.2 MeV and the tail at 6.5 MeV.

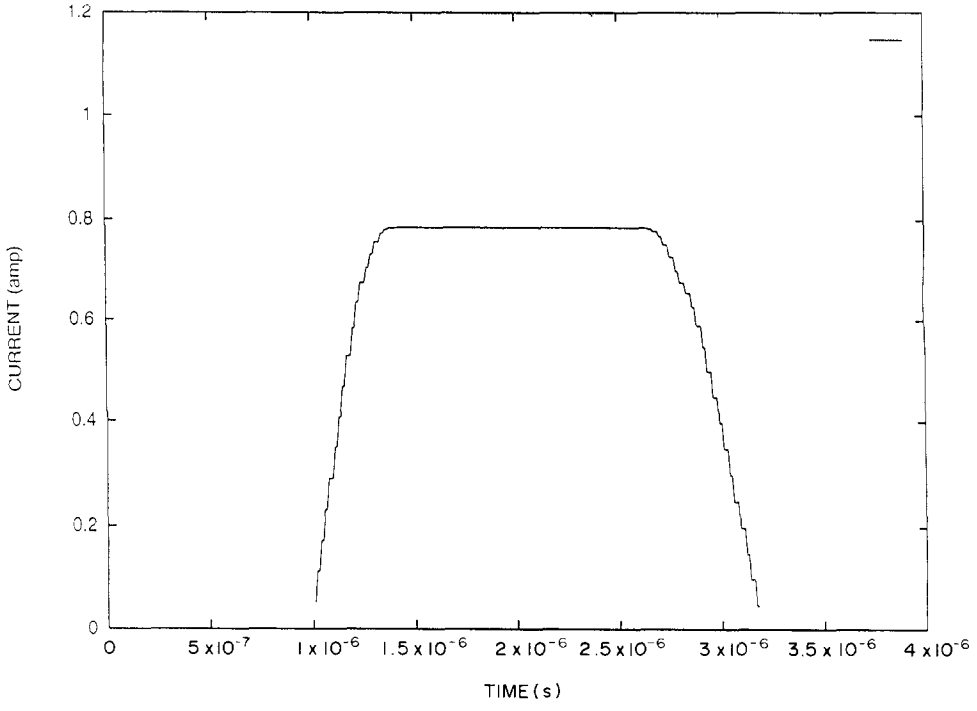

Fig. 3. Pulse acceleration at constant current to 5.7 MeV.

6. Cost optimization

Although the physics of the Elise design is the main topic of this paper, the design has also been optimized, so that the performance of the machine (measured in total pulse energy) has been maximized subject to the constraints of a maximum allowable cost, as well as subject to the requirement of meeting the goals of the physics experiments. To obtain a design, an approximate cost algorithm was developed that included relevant engineering constraints and data. A MATHEMATICA-based design code was written (in addition to an Excel spreadsheet code), which then laid down each half-lattice period, determining core and quadrupole lengths, evaluating beam parameters and calculating incremental costs for each of these components. Cost optimizations were then made that scanned over parameters such as length and width of the amorphous iron cores, ion energy and pulse duration.

6.1. Design algorithm

We now outline the basic algorithms used in the code. Many parameters of the machine and their justification have already been mentioned. The ion mass (39.1 a.m.u.) and charge state (+1), injected beam energy (2.0 MeV) and normalized emittance (1.0×10^{-6} m rad), initial current (0.78 A), pipe radius for each beam (2.33 cm) and quadrupole voltage (118 kV) are set by previous considerations.

Each induction core is an annular cylinder of ferromagnetic material. The beam acts similarly to the secondary of a one-turn transformer as it threads the core. The core allows a voltage increment V_c for a duration Δt_V than can be shown by use of Faraday's law to satisfy

$$V_c \Delta t_V = \Delta B A_c f_p$$

Here, ΔB is the magnetic flux swing in the ferromagnetic material, f_p is the packing fraction of the magnetic material about (0.75), and $A_c = (R_0 - R_i)\, L_c$ is the area of the core. The ferromagnetic

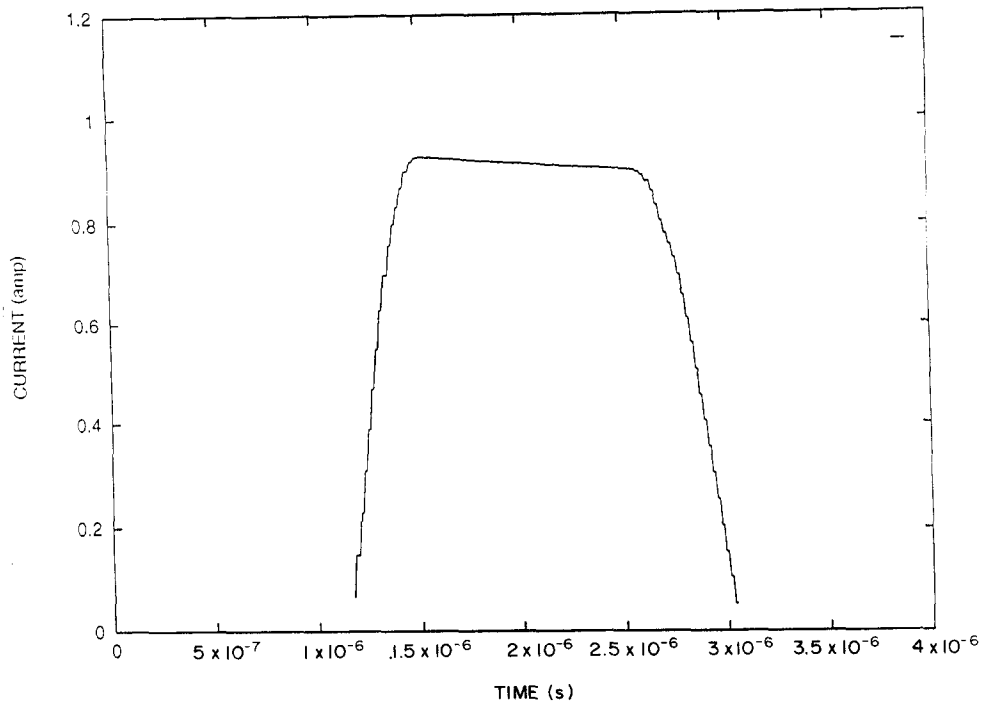

Fig. 4. Pulse acceleration with compression in time to 5.2 MeV. Note energy tilt.

material that we have chosen for cost and efficiency considerations is Metglas®. Note that Δt_V is the full width at half-maximum (FWHM) of the voltage pulse and that, for a flat-top current pulse of duration Δt_I, the current rise and fall each last about 0.3 μs, so that current flat-top of 1.5 μs consumes about 2.5 μs in effective voltage pulse.

Several cores of differing inner and outer radius can be combined on a common mandrel to form a single acceleration 'cell', and several cells may be stacked longitudinally to form a 'module', enclosed in a metallic can. Present designs generally have one module per half-lattice period. One or two core lengths are chosen to allow flexibility in adding cells. This flexibility is needed because space must be reserved at the end of each module (5.1 cm) to allow room for high voltage feedthroughs to power ESQs, and the location of these feedthroughs generally cannot be too close to the acceleration gaps between quadrupole, or too close to the groundplate (approximately midway between the gaps). By combining different core lengths, a large fraction of the longitudinal space can be utilized, avoiding excluded areas and maximizing the number of volt seconds per meter available for acceleration and pulse length. The algorithm used in the MATHEMATICA code at each half-lattice period adds, from a selection of two cell types, the maximum core area that can be added within the constraint imposed by the length of the half-lattice period and the constraint that the module end does not lie too close to the acceleration gap.

The voltage on each core has an upper limit determined primarily by the voltage stand-off capabilities of the thyratron switch, as used for initiation of the discharge of the pulse that forms network, which provides the energy storage and pulse shaping of the voltage pulse. If the core is chosen so that its area results in a volt second capability that is smaller than the product of the maximum voltage times the required pulse length for the current pulse, then either the voltage must be less than the maximum voltage, or the head and tail voltages must differ. The algorithm used in the optimization code adds the maximum

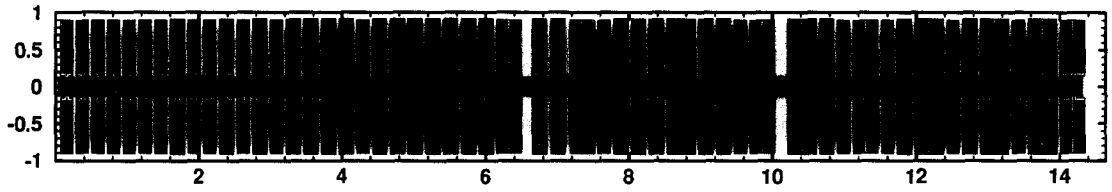

Fig. 5. An optimal Elise configuration of cores and quadrupoles.

voltage increment to the tail of the pulse and uses whatever voltage is permissible to the head, assuming a trapezoidal pulse shape. This gives a larger velocity to the tail than to the head, which in turn causes the pulse duration to decrease—a technique that will be used in fusion drivers to compress the pulse length as the beam accelerates. In most of the optimization runs, however, the machine is designed using a constant-current acceleration schedule, in which the pulse duration is constant.

Enough flexibility exists in the designs so that the schedules described in Section 5 can also be carried out. Once the energy increment is chosen, the length of the half-lattice period is determined by ensuring that the envelope of the beam remains 'matched', i.e.

$$\eta E' L / T^{1/2} = \text{constant}$$

Here, η is the longitudinal occupancy of the quadrupole relative to the half-lattice period, E' is the transverse electrostatic gradient in the quadrupole, L is the half-lattice period and T is the particle kinetic energy. This condition is applied at the beginning of each quadrupole, after acceleration from the previous acceleration gap. This leads to a half-lattice period that increases approximately as the square root of the energy when the quadrupole voltage is fixed.

With the energy increment and half-lattice period length determined, the analytic formulas for the beam properties at the head and tail of the pulse, such as the average and maximum beam radius, and depressed and undepressed tunes may be calculated using the results of Section 4. The current and line charge density are calculated at the head and the tail, ignoring space-charge effects and assuming charge conservation. These assumptions lead to the current at the head of the beam that satisfies

$$I_{h_i} \cong \left[\frac{I_h}{1 - (L/v_h^2)(\mathrm{d}v_h/\mathrm{d}t)} \right]_{i-1}$$

Here I_h, V_h, and $\mathrm{d}V_h/\mathrm{d}t$ are the current, velocity and velocity time derivative, evaluated at the head of the pulse, and i indicates the number of the acceleration gap. A similar expression is used for the tail of the pulse. A uniform velocity time derivative is assumed across the pulse $[\mathrm{d}v/\mathrm{d}t \cong (v_t - v_h)/(t_t - t_h)]$, where subscript t indicates tail, and t is arrival time at the end of a particular acceleration gap. This formula is kinematic only, since space charge has been neglected. However, for the purposes of optimization and initial design scoping, the approximations are adequate.

More accurate fluid and particle codes (cf. Sections 3–5 are used for beam studies. Simple models for the quadrupole dimensions, as well as the cell and module dimensions, are incorporated into the code and a graphical display of the entire machine is then generated (see Fig. 5).

6.2. Costing algorithms

Algorithms for costing the three major variable costs have been incorporated into the code. These are transport, acceleration modules and pulsers. A simplified cost model for the quadrupole transport can be expresses as

$$C_{\text{transport}} = C_{\text{fixed}} + C_{\text{hlp}} N_{\text{hlp}}$$

Here, $C_{\text{transport}}$ is the total transport cost in US dollars, and is composed of C_{fixed}, the fixed cost for developing the quad system (estimated to be about \$180 000 for a machine such as Elise), plus the cost per quad C_{hlp} times the number of quads (equal to the number of half-lattice periods N_{hlp}). C_{hlp} has been estimated to be \$25 000.

The second major cost component is the cost of the acceleration modules C_{modules}, which can be written in the form

$$C_{\text{modules}} = C_{\text{mod0}} N_{\text{mod}} + C_{\text{mod1}} (R_{\text{mo}}/1\ \text{m})^2 N_{\text{mod}} + C_{\text{core}} N_{\text{core}} + C_{\text{Met}} \rho_{\text{Met}} N_{\text{cell}} U_{\text{cell}}$$

Here, C_{mod0} ($\cong$ \$3900) is the cost associated per module, apart from costs that scale with outer module radius R_{mo}; C_{mod1} ($\cong$ \$8900) is the cost associated per module that scales with R_{mo}^2; C_{core} ($\cong$ \$2400) is the cost associated with each individual core; and C_{Met} ($\cong$ \$6–14 kg^{-1}) is the cost per kilogram of Metglas®, which includes mechanical costs that scale with accelerator weight (about \$3 kg^{-1}), with winding the cores (about \$3 kg^{-1}), and the price of Metglas® (from \$3.3 kg^{-1} to \$8.4 kg^{-1}, depending on the core length and variety of Metglas®). The variables N_{core}, N_{cell} and N_{mod} represent the total number of cores, cells and modules, respectively, ρ_{Met} is the density of Metglas® ($\cong$ 7180 kg m^{-3}) and U_{cell} is the volume of Metglas® in each accelerating cell.

The third major cost component is the cost of pulse power C_{pulser}, which has two major parts. One part arises because of the power limitations of the thyratron switches (about 72 MW) and the associated costs per pulser, and the second part is the marginal cost per joule of stored energy. These can be summarized approximately as

$$C_{\text{pulser}} \cong \$\ N_{\text{cell}} U_{\text{cell}} L [C_{\text{pulW}} + C_{\text{pulJ}} (1.25/\Delta T_V)]$$

Here, $C_{\text{pulW}} \cong 0.17$ kW^{-1} is the cost per kilowatt of a 72 MW pulser, and $C_{\text{pulJ}} \cong$ \$50 J^{-1} is the cost per joule of core energy loss, which requires energy storage in the pulse-forming line that makes up the pulser. L is the energy loss per cubic meter of core volume associated with eddy currents and hysteresis losses, and is given approximately by

$$L \cong (400\ \text{J m}^{-3}) \left(\frac{\Delta B}{2T}\right)^2 \left(\frac{2\ \mu\text{s}}{\Delta t_V}\right)$$

The power requirements on the pulser have been assumed to be 25% larger than the energy loss in the core L divided by the effective pulse length Δt_V, because of other energy loss mechanisms, such as those that occur in the pulse-forming line.

The total direct cost is taken to be the sum of $C_{\text{transport}}$, C_{modules} and C_{pulser}.

6.3. Results from the cost optimization

The results of the cost study yielded insight in four areas of optimization: core length, module outer radius, pulse duration and final ion energy. Fig. 6 illustrates a series of models for different core lengths. The dollars per joule value is plotted as a function of the number of cores per cell. The pulse duration is held constant in this series of curves. On a given curve, increasing the number of cores per cell increases the outer radius of the induction modules. As the core length increases, since the volt seconds per core and, hence, core area are fixed, the outer radius decreases at a fixed number of cores per cell. Two observations can be made. The outer radius of the machine is optimum at about 1.0 m, while the optimum tape width occurs at 14.2 or 17.0 cm (or a combination of the two). An optimum in the module radius occurs because the core volume and, hence, core cost increase proportionally to radius, which results in slow acceleration and, hence, in high transport costs. Also, when the number of cores per cell is too small, the weight of the Metglas® becomes smaller, but the total length of the machine increases and the cost is dominated by transport costs. The optimum in length occurs because using longer cores reduces the amount of unused spaces in between cells, increasing the long itudinal packing fraction of core material from about 0.5 for 5.1 cm cores to over 0.6 for the longer cores. Additionally, the longer cores have

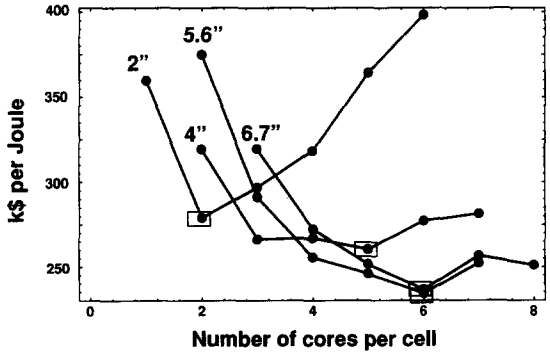

Fig. 6. Costs vs. core lengths for Elise.

a substantially lower unit cost per kilogram (about \$3.30 kg^{-1} vs. about \$8.4 kg^{-1} for the 5.1 cm cores).

Another result of the cost optimization is the trade-off between pulse duration and ion energy, when the cost of the machine is constrained. In this trade-off, it was found that, if the goal is to design a machine in which the pulse energy is maxmized, then increasing the pulse duration tends to win over increasing ion energy. Increasing the pulse duration requires larger and more costly cores but, because the core losses are reduced per unit core volume (because dB/dt is reduced) and the power requirements per unit core volume (proportional to $1/\Delta t_V$) are reduced, the pulser requirements increase less rapidly with increasing pulse duration than with ion energy. Therefore, for a fixed cost but maximizing pulse energy gained in the accelerator, a longer pulse duration is preferred. However, because Elise is being designed to demonstrate a variety of acceleration schedules, as well as to deliver maximal pulse energy, a maximum flat-top current pulse duration of about 1.5 μs is now being considered, so that the bunch length of the beam remains a reasonably small fraction of the accelerator length, and 'load and fire' scenarios and bunch length compression schedules can be carried out.

7. Elise diagnostics

There will be three basic types of detector device for the measurement of the beam current and transverse beam distributions: current-measuring Faraday cups, beam slits or pinholes, and current-profile measuring harps. The Faraday cups are fairly conventional units, changing little from the types used on previous LBNL experiments, such as MBE-4 and SBTE. The emittance-measuring slit-harp combination will evolve from the present design of slits with cups to slits with multiple-detector arrays (harp-like constructions). Harp detectors will have wires 0.05 mm thick and, possibly, active amplifier arrays at the detector elements. They can be fabricated with the necessary precision using electroformed nickel elements.

At least two locations in the matching section design will allow the breaking of the vacuum system and longitudinal separation to insert intercepting diagnostics for the measurement of the transverse phase space in two planes, and the total current. The accelerator is about 14 m long, with only two non-accelerating gaps about 6.5 and 10 m downstream of the matching section. Double-slit position and emittance-measuring devices will be installed in the matching section and in a box that will move downstream as the accelerator is assembled. Some of the hardware is available from earlier experiments.

The two non-accelerating locations in the lattice are the only places where intercepting diagnostics can easily be inserted without breaking the vacuum. The idea here is to measure the beam profile in two directions with a retractable slit–Faraday cup combination, utilizing only the 2.0 cm gap between the quadrupole end-plates. To minimize the possibility of breakdown between the diagnostics and the quadrupole end-plates, which are held at a potential of 59–70 kV, the gap between those particular quadrupoles may be increased from 2 to 4 cm. The quadrupole voltages immediately upstream and downstream would need to be adjusted somewhat to account for the break in the FODO structure.

After the accelerator has been built, double-slit, or slit-harp phase space measuring devices will be inserted as needed, by creating longitudinal space by breaking the vacuum between quadrupoles and inserting the diagnostics apparatus. The accelerator is being designed so that this will take less than one day.

The accelerating voltage waveforms will be measured using a capacitive voltage divider inserted in each accelerating cell. These measurements are essential for determining the proper operation of the accelerating cells, as well as for input into the longitudinal dynamics codes for evaluation of the accelerator experiments.

Data for the momentum distribution along an ion bunch are essential for studying longitudinal beam dynamics. With MBE-4 and SBTE, these parameters were measured using an electrostatic spectrometer with an energy resolution of better than 0.5%. For Elise, this technique is more

difficult, because the beam energies are a factor of 5 or more higher. Preliminary designs have been started for a magnetic analyzer that would be usable to full beam energy. This device is rather large and expensive, and would probably be useful only as an end-of-accelerator diagnostic.

The investigation of an energy measurement technique that determines the ion energy by time of flight is in progress. At the point of measurement the Elise beam is collimated to a pencil beam and gated to a short pulse (tens of nanoseconds). Measurement of the flight time from this point to a downstream detector (Rogowski loop, if adequately sensitive, or else a Faraday cup) can yield measurements of the beam energy vs. time in the pulse. This technique has the advantages of permitting measurements at any of the diagnostic locations within Elise and should be much less expensive than a magnetic spectrometer.

Acknowledgements

This work was performed under the auspices of the US Department of Energy at the Lawrence Berkeley National Laboratory under Contract DE-AC03-76SF00098 and at Lawrence Livermore National Laboratory under contract W-7405-ENG-48.

References

[1] C.M. Celata et al., Nuovo Cimento 106A(11) (1993) 1631.
[2] J.W. Kwan, R.O. Bangerter, A. Faltens, C. Peters and L.L. Reginato, Elise plans and progress, Fus. Eng. Des. 32–33 (1996) 299.
[3] E.A. Abramyan and V.A. Gaponov, Atom. Energ. 20 (1966) 385.
[4] W.B. Herrmannsfeldt, EGUN–an electron optics and gun design program, SLAC-Rep. 331, 1988.
[5] A. Friedman, D. Grote and I. Haber, Three-dimensional particle simulation of heavy-ion fusion beams, Phys. Fluids B4 (1992) 2203.
[6] C.L. Chang et al., Three-dimensional modeling of accelerators, Los Alamos Rep. LA-11857-C, 1990.
[7] S. Yu et al., Driver-scale ion injector experiments.
[8] D.W. Hewett and D.J. Larsen, The best of Gymnos: a user's guide, LLNL Rep. UCRL-ID-110499, May 1992.
[9] E.P. Lee, T.J. Fessenden and L.J. Laslett, IEEE Trans. Nucl. Sci. NS-32 (1985) 2489. E.P. Lee, The beam envelope equation—systematic solution for a periodic quadrupole lattice with space charge, Rep. LBL-37050, April 1995.
[10] C.H. Kim and L. Smith, Particle Accel. 85 (1986) 101.
[11] E. Henestroza, HIFAR Note 379, LBNL, March 1989.
[12] W.M. Sharp, Proc. 1995 Particle Accelerator Conf., Dallas, TX, 1–5 May 1995.